

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently amended) A method of implementing the address resolution protocol (ARP) in a computing platform having a plurality of processors interconnected by a non-Ethernet physical network, comprising:

defining a topology of an Ethernet network to be emulated on the computing platform, the topology including processor nodes and a switch node;

assigning a set of processors from the plurality to be processors to act as the processor nodes;

assigning a processor to act as the switch node and to emulate an Ethernet switch;

allocating virtual interfaces for the underlying non-Ethernet physical network, the virtual interfaces providing direct software communication paths between two processors connected to the non-Ethernet physical network, wherein each virtual interface has a corresponding identification;

a first processor node, in response to needing to communicate an IP packet to a target IP node for which the first processor node has an IP address but insufficient, corresponding lower layer address information, communicating an ARP request to the switch node via the non-Ethernet physical network, wherein the ARP request includes an IP address for the target node;

the switch node communicating, via the non-Ethernet physical network, the ARP request to all other processor nodes in the emulated Ethernet network;

a second processor node that is assigned the IP address for the target IP node receiving the ARP request and issuing an ARP reply in response;

the second processor node having its ARP table programmed to associate the IP address of the first node with a corresponding virtual interface for the underlying non-Ethernet physical network;

the first processor receiving and processing the ARP reply;

the first processor node having its ARP table programmed to associate the IP address of the target IP node with a corresponding virtual interface for the underlying non-Ethernet physical network;

wherein for subsequent unicast IP communication between the first and second nodes, the first and second processor nodes respectively use their ARP tables and the virtual interfaces associated therewith to communicate directly between processor nodes over the non-Ethernet physical network, avoiding the switch node; and wherein a subset of the processors are organized as a cluster and wherein one of the processors in the cluster is a load balancing processor node, and wherein, when any processor in the cluster issues an ARP request, the ARP request is programmed to include a MAC address for the load balancing processor node rather than for the requesting node, so that the reply to the ARP request is sent to the load balancing processor node rather than to the requesting node, and when the load balancing node receives the ARP reply it distributes said reply to all nodes in the cluster.

2. (Original) The method of claim 1 wherein the underlying physical network is a point-to-point mesh connecting the plurality of processors.
3. (Canceled)
4. (Original) The method of claim 1 wherein the switch node is in communication with an external IP network, and wherein the act of communicating an ARP reply includes identifying that the ARP reply is from a processor node in the platform.
5. (Currently amended) An address resolution protocol (ARP) system, comprising:
a computing platform having a plurality of processors connected by an underlying non-Ethernet physical network;
logic, executable on one of the processors, to define a topology of an Ethernet network to be emulated on the computing platform, the topology including processor nodes and a switch node;
logic, executable on one of the processors, to assign a set of processors from the plurality to be processors to act as the processor nodes;

logic, executable on one of the processors, to allocate virtual interfaces for the underlying non-Ethernet physical network to provide direct software communication paths between two processors connected to the non-Ethernet physical network, wherein each virtual interface has a corresponding identification;

each processor node having ARP request logic to communicate an ARP request to the switch node, wherein the ARP request includes an IP address;

the switch node including ARP request broadcast logic to communicate via the non-Ethernet physical network the ARP request to all other processor nodes in the emulated Ethernet network;

each processor node having ARP reply logic to determine whether it is the processor node associated with the IP address in an ARP request and, if so, to issue an ARP reply, and having logic to program its ARP table to associate the IP address of the ARP requester with a corresponding virtual interface for the underlying non-Ethernet physical network;

each processor node further having logic, responsive to ARP replies, to program its ARP table to associate the IP address of the ARP replier with a corresponding virtual interface for the underlying non-Ethernet physical network;

wherein for subsequent unicast IP communication between the first and second nodes, the first and second processor nodes respectively use their ARP tables and the virtual interfaces associated therewith to communicate directly between processor nodes over the non-Ethernet physical network, avoiding the switch node; and wherein a subset of the processors are organized as a cluster and wherein one of the processors in the cluster is a load balancing processor node, and wherein, when any processor in the cluster issues an ARP request, the ARP request is programmed to include a MAC address for the load balancing processor node rather than for the requesting node, so that the reply to the ARP request is sent to the load balancing processor node rather than to the requesting node, and when the load balancing node receives the ARP reply it distributes said reply to all nodes in the cluster.

6. (Original) The system of claim 5 wherein the underlying physical network is a point-to-point mesh connecting the plurality of processors.

7. (Canceled)
8. (Original) The system of claim 5 wherein the switch node is in communication with an external IP network, and wherein the processor node ARP reply logic includes logic to identify that the ARP reply is from a processor node in the platform.
9. (Previously presented) The method of claim 1 wherein the second processor node includes driver logic to modify the ARP request to include virtual interface information so that the ARP table of the second processor node associates the IP address of the first node with virtual interface information for the first processor node.
10. (Previously presented) The method of claim 1 wherein the first processor node includes driver logic to modify the ARP reply to include virtual interface information so that the ARP table of the first processor node associates the IP address of the second node with virtual interface information for the second processor node.
11. (New) A method of implementing the address resolution protocol (ARP) in a computing platform having a plurality of processors interconnected by a non-Ethernet physical network, comprising:
defining a topology of an Ethernet network to be emulated on the computing platform, the topology including processor nodes and a switch node;
assigning a set of processors from the plurality to be processors to act as the processor nodes;
assigning a processor to act as the switch node and to emulate an Ethernet switch so that the switch nodes and processor nodes form an emulated Ethernet network;
allocating virtual interfaces for the underlying non-Ethernet physical network, the virtual interfaces providing direct software communication paths between two processors connected to the non-Ethernet physical network, wherein each processor node has corresponding virtual interface address information and said corresponding virtual interface address information is specified in MAC address format and has a predefined setting of select address bits to identify the MAC address as carrying virtual interface address information within the MAC address format;

a first processor node, in response to needing to communicate an IP packet to a target IP node for which the first processor node has an IP address but insufficient, corresponding lower layer address information, communicating an ARP request to the switch node via the non-Ethernet physical network, wherein the ARP request includes an IP address for the target node;

the switch node communicating, via the non-Ethernet physical network, the ARP request to all other processor nodes in the emulated Ethernet network;

a second processor node that is assigned the IP address for the target IP node receiving the ARP request and issuing an ARP reply in response;

the second processor node having its ARP table programmed to associate the IP address of the first node with its corresponding virtual interface address information specified in MAC address format;

the first processor receiving and processing the ARP reply;

the first processor node having its ARP table programmed to associate the IP address of the target IP node with its corresponding virtual interface address information specified in MAC address format.

12. (New) An address resolution protocol (ARP) system, comprising:

a computing platform having a plurality of processors connected by an underlying non-Ethernet physical network;

logic, executable on one of the processors, to define a topology of an Ethernet network to be emulated on the computing platform, the topology including processor nodes and a switch node;

logic, executable on one of the processors, to assign a set of processors from the plurality to be processors to act as the processor nodes;

logic, executable on one of the processors, to allocate virtual interfaces for the underlying non-Ethernet physical network to provide direct software communication paths between two processors connected to the non-Ethernet physical network, wherein each processor node has corresponding virtual interface address information and said corresponding virtual interface address information is specified in MAC address format and has a predefined setting of select address bits to identify the

MAC address as carrying virtual interface address information within the MAC address format;
each processor node having ARP request logic to communicate an ARP request to the switch node, wherein the ARP request includes an IP address; the switch node including ARP request broadcast logic to communicate via the non-Ethernet physical network the ARP request to all other processor nodes in the emulated Ethernet network;
each processor node having ARP reply logic to determine whether it is the processor node associated with the IP address in an ARP request and, if so, to issue an ARP reply, and having logic to program its ARP table to associate the IP address of the ARP requester with its corresponding virtual interface address information specified in MAC address format;
each processor node further having logic, responsive to ARP replies, to program its ARP table to associate the IP address of the ARP replier with its corresponding virtual interface address information specified in MAC address format;
wherein for subsequent unicast IP communication between the first and second nodes, the first and second processor nodes respectively use their ARP tables and the virtual interfaces associated therewith to communicate directly between processor nodes over the non-Ethernet physical network, avoiding the switch node.